Applying Maxi-adjustment to Adaptive Information Filtering Agents

Raymond Lau and **Arthur H.M. ter Hofstede**
Cooperative Information Systems Research Centre
Queensland University of Technology
GPO Box 2434, Brisbane Qld 4001, Australia
E-mail:{raymond, arthur}@icis.qut.edu.au

Peter D. Bruza
Distributed Systems Technology Centre
The University of Queensland
Brisbane, Qld 4072, Australia
E-mail:bruza@dstc.edu.au

arXiv:cs/0003014v2 [cs.AI] 8 Mar 2000

Abstract

Learning and adaptation is a fundamental property of intelligent agents. In the context of adaptive information filtering, a filtering agent's beliefs about a user's information needs have to be revised regularly with reference to the user's most current information preferences. This learning and adaptation process is essential for maintaining the agent's filtering performance. The AGM belief revision paradigm provides a rigorous foundation for modelling rational and minimal changes to an agent's beliefs. In particular, the maxi-adjustment method, which follows the AGM rationale of belief change, offers a sound and robust computational mechanism to develop adaptive agents so that *learning autonomy* of these agents can be enhanced. This paper describes how the maxi-adjustment method is applied to develop the learning components of adaptive information filtering agents, and discusses possible difficulties of applying such a framework to these agents.

Introduction

With the explosive growth of the Internet and the World Wide Web (*Web*), it is becoming increasingly difficult for users to retrieve relevant information. This is the so-called problem of *information overload* on the Internet. Augmenting existing Internet search tools with personalised information filtering agents is one possible method to alleviate this problem. Adaptive information filtering agents are computer systems situated on the Web. They *autonomously* filter the incoming stream of information on behalf of the users. Users' information needs will change over time, and information filtering agents must be able to revise their beliefs about the users' information needs so that the accuracy of the filtering process can be maintained. The **AGM** belief revision paradigm (Alchourrón, Gärdenfors, & Makinson 1985) provides a rich and rigorous foundation for modelling such revision processes. It enables an agent to modify its beliefs in a rational and minimal way. **Maxi-adjustment** (Williams 1996; Williams 1997a) is a specific change strategy that follows the AGM's rationale of belief revision. In particular, it transmutes the underlying entrenchment ranking of beliefs in an absolute minimal way under maximal information inertia. In information retrieval models (Salton & McGill 1983; Salton 1989), information objects are often assumed independent unless semantic relationships among them can be derived. This intuition coincides with the underlying assumption of the maxi-adjustment strategy. The advantage of employing the maxi-adjustment strategy as the agents' learning mechanism is that semantic relationships among information items can be taken into account during the agents' learning and adaptation processes. Less users' relevance feedback (Salton & Buckley 1990) is required to train the filtering agents, and hence a higher level of *learning autonomy* can be achieved when compared with the other learning approaches employed in adaptive information agents (Billsus & Pazzani 1999; Moukas & Maes 1998; Balabanovic 1997; Pazzani, Muramatsu, & Billsus 1996; Armstrong *et al.* 1995). This paper focuses on the application of the maxi-adjustment method to the development of learning mechanisms in adaptive information filtering agents. Moreover, difficulties of applying such a framework to the filtering agents are discussed.

The Adaptive Filtering Agent

Figure 1 is an overview of the major functional components of an adaptive information filtering agent. The focus of this paper is on the *learning* component of the *adaptive information filtering agent*. The filtering agent's memory holds the current representation of a user's information needs. In particular, the notion of belief state (Gärdenfors 1988) is used to represent these information needs. The learning component accepts a user's relevance feedback about filtered Web documents. Based on the proposed induction method, these feedback is converted to the beliefs of a user's information needs. The maxi-adjustment strategy is then applied to revise the beliefs stored in the agent's memory with respect to these newly induced beliefs. As the user's information needs will change over time, such a belief revision process needs to be conducted repeatedly. Technically, the learning behavior demonstrated by the agent is a kind of *reinforcement learning* (Langley 1996). The matching component filters out relevant

information from a stream of incoming Web documents. It is underpinned by logical deduction. In other words, if the representation of a Web document is logically entailed by the set of beliefs stored in the agent's memory, the Web document will be considered as relevant and presented to the user. According to the user's relevance feedback, new beliefs may be added to, or existing beliefs may be contracted from the agent's memory. The adaptive information filtering agent is one of the main elements in the agent-based information filtering system (AIFS) (Lau, Hofstede, & Bruza 1999).

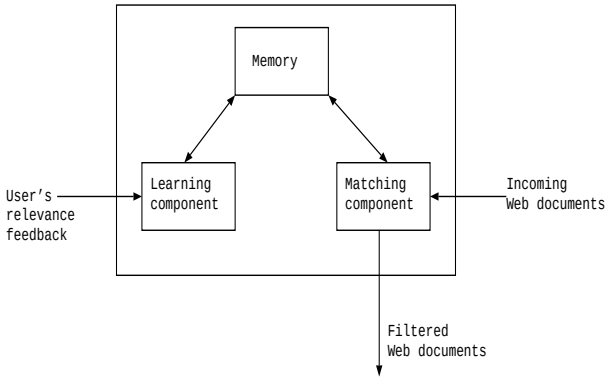

Figure 1: An Overview of Adaptive Information Filtering Agent

The AGM Belief Revision Paradigm

The AGM paradigm (Alchourrón, Gärdenfors, & Makinson 1985) provides a rigorous foundation for modelling consistent and minimal changes to an agent's beliefs. In particular, belief revision is taken as transitions among belief states (Gärdenfors 1988). A belief state can be represented by a belief set K which is a theory in a propositional language $\mathcal{L}$ (Gärdenfors & Makinson 1988). For the discussion in this paper, it is assumed that $\mathcal{L}$ is the classical propositional language. In the AGM framework (Alchourrón, Gärdenfors, & Makinson 1985; Gärdenfors 1988; Gärdenfors 1992; Gärdenfors & Makinson 1988), three principle types of belief state transitions are identified and modelled by corresponding belief functions: *expansion* (K^+_α), *contraction* (K^-_α), and *revision* (K^*_α). The process of belief revision can be derived from the process of belief contraction and vice versa through the so-called *Levi Identity* i.e. $K^*_\alpha = (K^-_{\neg\alpha})^+_\alpha$, and the *Harper Identity* i.e. $K^-_\alpha = K \cap K^*_{\neg\alpha}$. Essentially, the AGM framework proposes sets of postulates to characterise these functions such that they adhere to the rationales of *consistent* and *minimal* changes. In addition, it also describes the constructions of these functions based on various mechanisms. One of them is *epistemic entrenchment* ($\leqslant$) (Gärdenfors & Makinson 1988). For instance, if α, β are beliefs in a belief set K, $\alpha \leqslant \beta$ means that β is at least as entrenched as α. If inconsistency arises after applying changes to a belief set, beliefs with the lowest degree of epistemic entrenchment are given up. Technically, epistemic entrenchment (Gärdenfors & Makinson 1988; Gärdenfors 1992) is a total preorder of the sentences (e.g. α, β, γ) in $\mathcal{L}$, and is characterised by the following postulates: **(EE1)**: If $\alpha \leqslant \beta$ and $\beta \leqslant \gamma$, then $\alpha \leqslant \gamma$; **(EE2)**: If $\alpha \vdash \beta$, then $\alpha \leqslant \beta$; **(EE3)**: For any α and β, $\alpha \leqslant \alpha \wedge \beta$ or $\beta \leqslant \alpha \wedge \beta$; **(EE4)**: When $K \neq K_\perp$, $\alpha \notin K$ iff $\alpha \leqslant \beta$, for all β; **(EE5)**: If $\beta \leqslant \alpha$ for all β, then $\vdash \alpha$.

It has been proved that an unique contraction function can be defined by the underlying epistemic entrenchment through the (C-) condition (Gärdenfors & Makinson 1988):

$$\text{(C-)} \qquad K^-_\alpha = \begin{cases} \{\beta \in K : \alpha < \alpha \vee \beta\} & \text{if } \nvdash \alpha \\ K & \text{otherwise} \end{cases}$$

where $<$ is the strict part of epistemic entrenchment defined above. Moreover, the (C-R) condition (Gärdenfors 1992) also ensures that if an ordering of beliefs satisfies (EE1)-(EE5), the contraction function, uniquely determined by (C-R), satisfies all but the *recovery* postulates for contraction.

$$\text{(C-R)} \qquad K^-_\alpha = \begin{cases} \{\beta \in K : \alpha < \beta\} & \text{if } \nvdash \alpha \\ K & \text{otherwise} \end{cases}$$

Nevertheless, for a computer based implementation, a finite representation of epistemic entrenchment ordering and a policy of iterated belief changes are required. Williams (Williams 1995; Williams 1997b) proposed the *finite partial entrenchment ranking* ($\mathbf{B}$) that ranked the sentences of a theory in $\mathcal{L}$ with the minimum possible degree of entrenchment ($\leqslant_{\mathbf{B}}$). Moreover, *maxi-adjustment* (Williams 1996; Williams 1997a) was proposed to transmute a finite partial ranking using an absolute measure of minimal change under maximal information inertia. Belief revision is not just taken as adding or deleting a sentence from a theory but the transmutation of the underlying entrenchment ranking. Williams (Williams 1995; Williams 1996; Williams 1997a) formally defined the following definitions for a computational model of belief revision.

Definition 1 *A finite partial entrenchment ranking is a function* $\mathbf{B}$ *that maps a finite subset of sentences in* $\mathcal{L}$ *into the interval* $[0, \mathcal{O}]$ *such that the following conditions are satisfied for all* $\alpha \in dom(\mathbf{B})$*:*

(PER1) $\{\beta \in dom(\mathbf{B}) : \mathbf{B}(\alpha) < \mathbf{B}(\beta)\} \nvdash \alpha$.
(PER2) If $\vdash \neg\alpha$ *then* $\mathbf{B}(\alpha) = 0$.
(PER3) $\mathbf{B}(\alpha) = \mathcal{O}$ *if and only if* $\vdash \alpha$.

The set of all partial entrenchment rankings is denoted $\mathcal{B}$. $\mathbf{B}(\alpha)$ is referred as the *degree of acceptance* of α. The explicit information content of $\mathbf{B} \in \mathcal{B}$ is $\{\alpha \in \text{dom}(\mathbf{B}) : \mathbf{B}(\alpha) > 0\}$, and is denoted $exp(\mathbf{B})$. Similarly, the implicit information content represented by $\mathbf{B} \in \mathcal{B}$ is $Cn(exp(\mathbf{B}))$, and is denoted $content(\mathbf{B})$. Cn is the classical consequence operator. In order to describe the epistemic entrenchment ordering ($\leqslant_{\mathbf{B}}$) generated from a finite partial entrenchment ranking $\mathbf{B}$, it is necessary to rank implicit sentences.

Definition 2 *Let α be a non tautological sentence. Let* $\mathbf{B}$ *be a finite partial entrenchment ranking. The degree of acceptance of α is defined as:*

$$degree(\mathbf{B},\alpha) = \begin{cases} \text{largest } j \text{ such that} \\ \{\beta \in exp(\mathbf{B}) : \mathbf{B}(\beta) \geq j\} \vdash \alpha \\ \qquad\qquad \text{if } \alpha \in content(\mathbf{B}) \\ 0 \qquad\qquad otherwise \end{cases}$$

The maxi-adjustment strategy transmutes a partial entrenchment ranking $\mathbf{B}$ based on the rationale of absolute minimal change under maximal information inertia. It is assumed that sentences in exp($\mathbf{B}$) (e.g. α, β) are independent unless logical dependence exists between them. In particular, α is defined as a *reason* of β if and only if $degree(\mathbf{B}, \alpha \rightarrow \beta) > \mathbf{B}(\beta)$.

Definition 3 *Let $\mathbf{B} \in \mathcal{B}$ be finite. The range of $\mathbf{B}$ is enumerated in ascending order as $j_0, j_1, j_2, \ldots, j_{\mathcal{O}}$. Let α be a contingent sentence, $j_m = degree(\mathbf{B}, \alpha)$ and $0 \leq i < \mathcal{O}$. Then the (α, i) maxi-adjustment of $\mathbf{B}$ is $\mathbf{B}^\star(\alpha, i)$ defined by:*

$$\mathbf{B}^\star(\alpha,i) = \begin{cases} (\mathbf{B}^-(\alpha,i)) & \text{if } i \leq j_m \\ (\mathbf{B}^-(\neg\alpha,0))^+(\alpha,i) & otherwise \end{cases}$$

where for all $\beta \in dom(\mathbf{B})$, $\mathbf{B}^-(\alpha, i)$ is defined as follows:

1. For β with $\mathbf{B}(\beta) > j_m$, $\mathbf{B}^-(\alpha,i)(\beta) = \mathbf{B}(\beta)$.

2. For β with $i < \mathbf{B}(\beta) \leq j_m$, assuming that $\mathbf{B}^-(\alpha,i)(\beta)$ for β is defined with $\mathbf{B}(\beta) \geq j_{m-k}$ for $k = -1, 0, 1, 2, \ldots, n-1$, then for β with $\mathbf{B}(\beta) = j_{m-n}$,

$$\mathbf{B}^-(\alpha,i)(\beta) = \begin{cases} i \quad \text{if } \alpha \vdash \beta \text{ or} \\ \quad \alpha \nvdash \beta \text{ and } \beta \in \Gamma \\ \quad \text{where } \Gamma \text{ is a minimal subset of} \\ \quad \{\gamma : \mathbf{B}(\gamma) = j_{m-n}\} \text{ such that} \\ \quad \{\gamma : \mathbf{B}^-(\alpha,i)(\gamma) > j_{m-n}\} \cup \Gamma \vdash \alpha \\ \mathbf{B}(\beta) \quad otherwise \end{cases}$$

3. For β with $\mathbf{B}(\beta) \leq i$, $\mathbf{B}^-(\alpha,i)(\beta) = \mathbf{B}(\beta)$.

For all $\beta \in dom(\mathbf{B}) \cup \{\alpha\}$, $\mathbf{B}^+(\alpha, i)$ is defined as follows:

$$\mathbf{B}^+(\alpha,i)(\beta) = \begin{cases} \mathbf{B}(\beta) \qquad \text{if } \mathbf{B}(\beta) > i \\ i \qquad\qquad \text{if } \alpha \equiv \beta \text{ or} \\ \qquad \mathbf{B}(\beta) \leq i < degree(\mathbf{B}, \alpha \rightarrow \beta) \\ degree(\mathbf{B}, \alpha \rightarrow \beta) \qquad otherwise \end{cases}$$

It has been stated that if $i > 0$ then $content(\mathbf{B}^\star(\alpha, i))$ is an AGM revision, and $content(\mathbf{B}^\star(\alpha, 0))$ satisfies all but the *recovery* postulates for AGM contraction (Williams 1996).

Knowledge Representation

A Web page is characterised by a set of weighted keywords based on traditional information retrieval (*IR*) techniques (Salton & McGill 1983; Salton 1989). At the symbolic level, each keyword k is mapped to the ground term of the *positive keyword* predicate *pkw* i.e. $pkw(k)$. Basically, $pkw(k)$ is a *proposition* since its interpretation is either true or false. The intended interpretation of these sentences is that they are satisfied in a document D i.e. $D \models pkw(k)$ if D is taken as a model (Chiaramella & Chevallet 1992; Lalmas 1998). For example, if d = { business, commerce, trade, ... } is the document representation at the keyword level, the corresponding representation at the symbolic level will be d = { $pkw(business)$, $pkw(commerce)$, $pkw(trade)$, ... }. Similarly, a user's information needs are also represented as a set of weighted keywords at the keyword level. However, this set of weighted keywords is derived from a set of relevant documents D^+ and a set of non-relevant documents D^- with respect to the user's information needs (Allan 1996; Buckley, Salton, & Allan 1994). Based on the frequencies of these keywords appearing in D^+ and D^-, it is possible to induce a preference ordering among the keywords with respect to the user's information needs. The basic idea is that a keyword appearing more frequently in D^+ is a more preferred keyword than another keyword that appears less frequently in D^+. Once this preference ordering is induced, it is taken as the epistemic entrenchment ordering of the corresponding beliefs. It is observed that the postulates of epistemic entrenchment is valid in the context of information retrieval in general (Lau, Hofstede, & Bruza 1999). For example, if $\{\alpha, \beta, \gamma\}$ is a set of information carriers (Bruza & Huibers 1994), $\alpha \leqslant \beta$ and $\beta \leqslant \gamma$ implies $\alpha \leqslant \gamma$. In other words, if an information searcher prefers information carrier γ rather than information carrier β, and information carrier β rather than information carrier α, they prefer retrieving γ rather than α. This characteristic of information carriers matches the epistemic entrenchment postulate e.g. (EE1) of beliefs.

Moreover, it is necessary to classify a keyword as *positive*, *neutral*, or *negative* (Kindo *et al.* 1997). Intuitively, positive keywords represent the information items in which the users interested. Negative keywords represent the information items that the users do not want to retrieve. Neutral keywords mean that these keywords are not useful for determining the users' interests. Eq.(1) is developed based on the *keyword classifier* (Kindo *et al.* 1997). It can be used to induce the preference value $pre(k)$ of a keyword k, and classify it as positive, negative, or neutral.

$$pre(k) = \epsilon \times \tanh\left(\frac{df(k)}{\xi}\right) \times \left(p(k_{rel}) \tanh\frac{p(k_{rel})}{p_{rel}} - (1 - p(k_{rel})) \tanh\frac{(1-p(k_{rel}))}{(1-p_{rel})}\right) \quad (1)$$

where ϵ is used to restrict the range of $pre(k)$ such that $-1 < pre(k) < 1$. The examples illustrated in this paper assume that $\epsilon = 0.9$. $df(k)$ is the sum of the number of relevant documents $df(k_{rel})$ and the number of non-relevant documents $df(k_{nrel})$ that contains the keyword k, and tanh is the hyperbolic tangent. The rarity parameter ξ is used to control rare or new keywords and is expressed as int($logN + 1$), where N is the total number of Web documents judged by a user, and int is an integer function that truncates the decimal values. $p(k_{rel})$ is the estimated probability that a docu-

ment containing keyword k is relevant and is expressed as the fraction $\frac{df(k_{rel})}{df(k_{rel})+df(k_{nrel})}$. p_{rel} is the estimated probability that a document is relevant. In our system, it is assumed that the probability that a Web document presented by the filtering agent and judged as relevant by a user is $p_{rel} = 0.5$. A positive value of $pre(k)$ implies that the associated keyword is positive, whereas a negative value of $pre(k)$ indicates a negative keyword. If $pre(k)$ is below a threshold value λ, the associated keyword is considered neutral. It is assumed that $\lambda = 0.5$ for the examples demonstrated in this paper. Basically a positive keyword k is mapped to $pkw(k)$, and a negative keyword k is mapped to $\neg pkw(k)$. There is no need to create the symbolic representations for neutral keywords. For $pkw(k)$ or $\neg pkw(k)$, the entrenchment rank $\mathbf{B}(\alpha_k)$ of the corresponding formula α_k is defined as:

$$\mathbf{B}(\alpha_k) = \begin{cases} |pre(k)| & \text{if } |pre(k)| \geq \lambda \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

The following is an example of computing the entrenchment rank $\mathbf{B}(\alpha_k)$ from a set of judged Web documents. It is assumed that there are a set of five documents having been judged as relevant (i.e. $|D^+| = 5$) and another set of five documents having been judged as non-relevant by a user (i.e. $|D^-| = 5$). Each document is characterised by a set of keywords e.g. $d_1 = \{business, commerce, system\}, d_2 = \{art, sculpture\}$. Table 1 summarises the frequencies of these keywords appearing in both D^+ and D^-, their preference values, and the entrenchment ranks of corresponding formulae.

Table 1: Representation of users' information preferences

Keywords	D^+	D^-	$pre(k)$	Formula:α_k	Rank:$\mathbf{B}(\alpha_k)$
business	5	0	0.856	$pkw(business)$	0.856
commerce	4	0	0.836	$pkw(commerce)$	0.836
system	2	2	0	-	-
art	0	5	-0.856	$\neg pkw(art)$	0.856
sculpture	0	3	-0.785	$\neg pkw(sculpture)$	0.785
insurance	1	0	0.401	-	-

Learning and Adaptation

Whenever a user provides relevance feedback for a presented Web document, the belief revision process can be invoked to learn the user's current information preferences. Conceptually, the filtering agent's learning and adaptation mechanism is characterised by the belief *revision* and *contraction* processes. For example, if $\Gamma = \{a_1, a_2, \ldots, a_n\}$ is a set of formulae representing a Web document $d \in D^+$, the belief revision process $(((K^*_{a_1})^*_{a_2})\ldots)^*_{a_n}$ is invoked for each $a_i \in \Gamma$, where K is the belief set stored in the filtering agent's memory. On the other hand, the belief contraction process $(((K^-_{a_1})^-_{a_2})\ldots)^-_{a_n}$ is applied for each $a_i \in \Gamma$ and $d \in D^-$. The sequence of revising or contracting the set of beliefs $a_1, a_2, \ldots, a_n$ is determined by their entrenchment ranks and whether it is a revision or contraction operation. At the computational level, belief revision is actually taken as the adjustment of entrenchment ranking $\mathbf{B}$ in the *theory base* $exp(\mathbf{B})$. Particularly, *maxi-adjustment* (Williams 1996; Williams 1998) is employed by the learning component of the filtering agent to modify the ranking of its beliefs in an absolute minimal way under maximal information inertia. As the input to the maxi-adjustment algorithm consists of a sentence α and its entrenchment rank i, the procedure described in the Knowledge Representation Section is used to induce the new rank for each $\alpha \in \Gamma$. Moreover, for our implementation of the maxi-adjustment algorithm, the maximal ordinal $\mathcal{O}$ in the interval $[0, \mathcal{O}]$ is chosen as 1.

One advantage of a symbolic representation of the filtering agent's domain knowledge is that semantic relationships among keywords can be captured. For example, if the keywords *business* and *commerce* are taken as synonymous, this semantic relationship can be modelled as a formula $pkw(business) \equiv pkw(commerce)$ in the filtering agent's memory. Moreover, classification knowledge such as *sculpture* is a kind of *art* can also be used by specifying the rule such as $pkw(sculpture) \rightarrow pkw(art)$. It is believed that capturing the semantic relationships among keywords can improve the effectiveness of the matching process (Hunter 1995; Nie, Brisebois, & Lepage 1995). In fact, by employing maxi-adjustment as the filtering agent's learning mechanism, these semantic relationships can be reasoned about during the reinforcement learning process. This could lead to a higher level of *learning autonomy* since changes to related keywords can automatically be inferred by the filtering agent. As a result, less users' relevance feedback may be required. The following examples assume that the formulae $pkw(business) \equiv pkw(commerce)$ and $pkw(sculpture) \rightarrow pkw(art)$ have been manually added to the the agent's memory through a knowledge engineering process.

Example 1:

The first example shows how adding one belief to the agent's memory will automatically raise the entrenchment rank of another related belief. It is assumed that the belief $\mathbf{B}(\neg pkw(sculpture)) = 0.785$ and the belief $\mathbf{B}(pkw(business)) = 0.856$ have been learnt by the filtering agent. If several Web documents characterised by the keyword *art* are also judged as non-relevant by the user later on, the preference value of the keyword *art* can be induced according to Eq.(1). Assuming that $pre(art) = -0.856$, the corresponding entrenchment rank can be computed as $\mathbf{B}(\neg pkw(art)) = 0.856$ according to Eq.(2). By applying $\mathbf{B}^\star(\neg pkw(art), 0.856)$ to the theory base $exp(\mathbf{B})$, the before and after images of the agent's *explicit beliefs* (i.e. $exp(\mathbf{B})$) can be tabulated in Table 2. Based on the maxi-adjustment al-

gorithm, $\mathbf{B}^+(\alpha, i)(\beta) = i$ if $\mathbf{B}(\beta) \leq i < degree(\mathbf{B}, \alpha \rightarrow \beta)$.

$\because \mathbf{B}(\neg pkw(sculpture)) \leq 0.856 < degree(\mathbf{B}, \neg pkw(art) \rightarrow \neg pkw(sculpture))$
$\therefore \mathbf{B}^+(\neg pkw(art), 0.856)(\neg pkw(sculpture)) = 0.856$

The implicit belief $\neg pkw(art) \rightarrow \neg pkw(sculpture)$ in $content(\mathbf{B})$ is derived from the explicit belief $pkw(sculpture) \rightarrow pkw(art)$ in the theory base $exp(\mathbf{B})$, and its degree of acceptance is 1 according to Definition 2. As the belief $\neg pkw(art)$ implies the belief $\neg pkw(sculpture)$ and the agent believes in $\neg pkw(art)$, the belief $\neg pkw(sculpture)$ should be at least as entrenched as the belief $\neg pkw(art)$ according to (PER1) of Definition 1 or (EE2). In other words, whenever the agent believes that the user is not interested in *art* (i.e. $\neg pkw(art)$), it must be prepared to accept that the user is also not interested in *sculpture* at least to the degree of the former. The proposed learning and adaptation framework is more effective than other learning approaches that can not take into account the semantic relationships among information items. This example demonstrates the automatic revision of the agent's beliefs about related keywords given the relevance feedback for a particular keyword. Therefore, less users' relevance feedback may be required during reinforcement learning. Consequently, *learning autonomy* of the filtering agent can be enhanced.

Table 2: Raising related beliefs

Formula:α	$\mathbf{B}(\alpha)$ Before	$\mathbf{B}(\alpha)$ After
$pkw(business) \equiv pkw(commerce)$	1.000	1.000
$pkw(sculpture) \rightarrow pkw(art)$	1.000	1.000
$pkw(business)$	0.856	0.856
$\neg pkw(sculpture)$	0.785	**0.856**
$\neg pkw(art)$	0	**0.856**

Example 2:
The second example illustrates the belief contraction process. In particular, how the contraction of one belief will automatically remove another related belief from the agent's memory if there is a semantic relationship between the underlying keywords. Assuming that more Web documents characterised by the keyword *sculpture* are judged as relevant by the user at a later stage, the belief $\mathbf{B}(pkw(sculpture)) = 0.785$ could be induced. As $\mathbf{B}^\star(\alpha, i) = (\mathbf{B}^-(\neg\alpha, 0))^+(\alpha, i)$ if $i > j_m$, where $i = 0.785$ and $j_m = 0$ in this example, $\mathbf{B}^\star(pkw(sculpture), 0.785)$ leads to the contraction of the belief $\neg pkw(sculpture)$ from the theory base $exp(\mathbf{B})$. Moreover, $\mathbf{B}^-(\neg pkw(sculpture), 0)(\neg pkw(art)) = 0$ is computed because $\mathbf{B}(\neg pkw(art)) = 0.856 = j_{m-n}$ and the set $\{\gamma : \mathbf{B}^-(\neg pkw(sculpture), 0)(\gamma) > 0.856\} \cup \{\neg pkw(art)\} \vdash \neg pkw(sculpture)$ is obtained. The before and after images of the filtering agent's explicit beliefs are tabulated in Table 3.

Example 3:

Table 3: Contracting related beliefs

Formula:α	$\mathbf{B}(\alpha)$ Before	$\mathbf{B}(\alpha)$ After
$pkw(business) \equiv pkw(commerce)$	1.000	1.000
$pkw(sculpture) \rightarrow pkw(art)$	1.000	1.000
$pkw(business)$	0.856	0.856
$pkw(sculpture)$	0	**0.785**
$\neg pkw(sculpture)$	0.856	**0**
$\neg pkw(art)$	0.856	**0**

This example demonstrates how multiple sentences from the same Web document judged by a user can be contracted from the agent's memory. If some Web documents characterised by *sculpture* and *business* have recently been judged as non-relevant by the user, both the belief $\mathbf{B}(pkw(sculpture)) = 0$ and the belief $\mathbf{B}(pkw(business)) = 0$ could be induced. Since beliefs with minimal entrenchment rank (i.e. 0) are not supposed to be stored in the agent's memory, maxi-adjustement for these two beliefs is still required so that they can be removed from the agent's memory. Basically, the contraction process $\mathbf{B}^-(\alpha, i)$ is applied to both sentences. As opposed to the belief expansion process where the most entrenched sentence is applied to the agent's memory first, the least entrenched belief is first contracted from the agent's memory for belief contraction. Consequently, $\mathbf{B}^-(pkw(sculpture), 0)$ is invoked first. As there is not other logically related sentences in the theory base $exp(\mathbf{B})$, the belief $pkw(sculpture)$ is simply removed from the agent's memory. Similarly, the sentence $pkw(business)$ is also removed from $exp(\mathbf{B})$ by applying $\mathbf{B}^-(pkw(business), 0)$. The before and after images of the filtering agent's explicit beliefs are tabulated in Table 4.

Table 4: Contracting multiple beliefs in one cycle

Formula:α	$\mathbf{B}(\alpha)$ Before	$\mathbf{B}(\alpha)$ After
$pkw(business) \equiv pkw(commerce)$	1.000	1.000
$pkw(sculpture) \rightarrow pkw(art)$	1.000	1.000
$pkw(business)$	0.856	**0**
$pkw(sculpture)$	0.785	**0**

Filtering Web Documents

In our current framework, the matching function of the filtering agent is modelled as logical deduction. Moreover, a Web document is taken as the conjunction of a set of formulae (Chiaramella & Chevallet 1992; Hunter 1995). The following example illustrates the agent's deduction process with reference to previous examples. The resulting belief sets $K = Cn(exp(\mathbf{B}))$ from examples 1, 2, and 3 in the previous section are used to determine the relevance of the following three Web documents:

$$\phi = \{pkw(business) \wedge pkw(art)\}$$
$$\varphi = \{pkw(sculpture) \wedge pkw(art)\}$$
$$\psi = \{pkw(business) \wedge pkw(commerce)\}$$

The filtering agent's conclusions about the relevance of the Web documents are summarised as follows:

Time: (t1)

$$K = Cn(\{pkw(business) \equiv pkw(commerce), pkw(sculpture) \rightarrow pkw(art), pkw(business), \neg pkw(sculpture), \neg pkw(art)\})$$
$$\therefore \quad K \nvdash \phi, \quad K \nvdash \varphi, \quad K \vdash \psi$$

Time: (t2)

$$K = Cn(\{pkw(business) \equiv pkw(commerce), pkw(sculpture) \rightarrow pkw(art), pkw(business), pkw(sculpture)\})$$
$$\therefore \quad K \vdash \phi, \quad K \vdash \varphi, \quad K \vdash \psi$$

Time: (t3)

$$K = Cn(\{pkw(business) \equiv pkw(commerce), pkw(sculpture) \rightarrow pkw(art)\})$$
$$\therefore \quad K \nvdash \phi, \quad K \nvdash \varphi, \quad K \nvdash \psi$$

As can be seen tentative conclusion drawn at time $t1$ may not hold when new information e.g. $K^{-}_{pkw(business)}$ is processed by the agent at time $t3$. Strictly speaking, the deduction process of the agent should be described as $K \mathrel{|\!\sim} \psi$, where $\mathrel{|\!\sim}$ is a nonmonotonic inference relation because the inferred beliefs (i.e. conclusions) will not grow monotonically. It is not difficult to see that this $\mathrel{|\!\sim}$ should belong to the class of nonmonotonic inference called *expectation inference* ($\mathrel{|\!\sim}_{\leqslant}$) (Gärdenfors & Makinson 1994). The basic idea of expectation inference is that given a sentence α of a propositional language $\mathcal{L}$, if α and the subset of sentences in a belief set K that is consistent with α (i.e. $\{\alpha\} \cup \{\beta \in K : \neg\alpha < \beta\}$) can classically entail another sentence β, $\alpha \mathrel{|\!\sim}_{\leqslant} \beta$ can be deduced. With reference to our examples, since K classically entails or not entails ϕ, φ and ψ, and $K \cup K = K$, the definition of expectation inference can trivially be applied to describe the characteristics of the inference mechanism in the filtering agent. Therefore, $K \mathrel{|\!\sim}_{\leqslant} \phi$, where K is the filtering agent's belief set and ϕ is the logical representation of a Web document, represents an inference conducted by the adaptive filtering agent.

Discussion

An alternative approach for developing the filtering agent's learning mechanism is to apply the (C-) or the (C-R) condition to revise the agent's beliefs. It has been stated that both the (C-) and the (C-R) conditions can be used to construct the same class of belief revision functions that satisfy the AGM postulates for belief revision (Gärdenfors 1992). The following example illustrates how belief revision may be conducted based on the (C-R) condition. Assuming that the filtering agent's initial belief set K is as follows:

$$K = \{pkw(business), pkw(commerce), pkw(sculpture), pkw(art)\}$$

Moreover,

$$pkw(business) \geqslant pkw(commerce) \geqslant pkw(sculpture) \geqslant pkw(art)$$

If a user perceives some Web documents characterised by the keyword *business* as non-relevant, the revision process $K^{*}_{\neg pkw(business)}$ will be invoked. Based on the Levi Identity, the sentence $pkw(business)$ should first be contracted from the belief set K. Consequently, all the beliefs from K will be contracted according to the (C-R) condition. The resulting belief set becomes $K^{'} = \{\neg pkw(business)\}$. Nevertheless, a user who does not require information objects about *business* may still be interested in information objects about *art* and *sculpture*. Therefore, applying (C-R) or (C-) to construct belief contraction or belief revision function seems producing drastic changes in the context of information retrieval and filtering. In fact, term independence is often assumed in information retrieval. This intuition is reflected in the vector space model of information retrieval (Salton & McGill 1983), where changing the weight of a particular keyword may not affect the others in the weight vector. Therefore, the maxi-adjustment strategy produces a better approximation in terms of revising or contracting beliefs about information objects with respect to a user's information needs.

Under the current framework, domain knowledge such as semantic relationships among information objects is transferred to the filtering agent's memory through a knowledge engineering process. Since not all semantic relationship is highly certain (i.e. assigning the maximal entrenchment rank), by applying belief revision to the agent's memory, the corresponding beliefs may be contracted from the memory over time. Therefore, domain knowledge perhaps needs to be transferred to the agent's memory periodically in accordance with the postulates of epistemic entrenchment. This can be taken as an off-line process to minimise its impact on the availability of the filtering agents. However, further investigation is required to apply such a background learning process to the filtering agents.

The belief set $K = content(\mathbf{B})$ is actually used by the agents to infer the relevance of Web documents. On the other hand, maxi-adjustment is employed to revise the theory base $exp(\mathbf{B})$ and to maintain its consistency after applying changes. Though maxi-adjustment ensures that the revised theory base is consistent, it is possible that the belief set K becomes inconsistent (i.e. $K \vdash \bot$) after applying changes such as $\mathbf{B}^{\star}(\alpha, i)$. The following is a classical example to explain such a problem. It is assumed that the set of explicit beliefs $exp(\mathbf{B})$ as well as their entrenchment ranking is as follows:

$$exp(\mathbf{B}) = \{\forall x[penguin(x) \rightarrow bird(x)],$$
$$\forall x[penguin(x) \rightarrow \neg fly(x)],$$
$$\forall x[bird(x) \rightarrow fly(x)]\}$$

$$\mathbf{B}(\forall x[penguin(x) \rightarrow bird(x)]) = 0.9,$$
$$\mathbf{B}(\forall x[penguin(x) \rightarrow \neg fly(x)]) = 0.7,$$
$$\mathbf{B}(\forall x[bird(x) \rightarrow fly(x)]) = 0.4$$

In the context of information retrieval and filtering, $penguin(x) \rightarrow bird(x)$ can be interpreted as: if a user is interested in information objects about *penguin*, it is likely that the user is also interested in information objects about *bird*. The other formulae in $exp(\mathbf{B})$ can be interpreted in similar way. Given the fact that the user is interested in *tweety* which is a penguin i.e. $penguin(tweety)$ and $\mathbf{B}(penguin(tweety)) = 0.8$, by applying maxi-adjustment $\mathbf{B}^{\star}(penguin(tweety), 0.8)$, the revised entrenchment ranking $new\mathbf{B}$ will be:

$$new\mathbf{B}(\forall x[penguin(x) \rightarrow bird(x)]) = 0.9,$$
$$new\mathbf{B}(penguin(tweety)) = 0.8,$$
$$new\mathbf{B}(\forall x[penguin(x) \rightarrow \neg fly(x)]) = 0.7,$$
$$new\mathbf{B}(\forall x[bird(x) \rightarrow fly(x)]) = 0.4$$

The degree of acceptance of implicit beliefs is computed based on definition 2:

$$degree(new\mathbf{B}, \neg fly(tweety)) = 0.7,$$
$$degree(new\mathbf{B}, fly(tweety)) = 0.4$$

As can be seen, even though $exp(new\mathbf{B}) \nvdash \bot$, it is clear that $content(new\mathbf{B}) \vdash \bot$, where $\vdash$ is classical derivability relation. If the filtering agent employs the belief set $K = content(new\mathbf{B})$ to deduce the relevance of Web documents, any documents will be considered as relevant. This problem must be addressed before the filter agents can be put to practical use. One possible solution is to make use of the degree of acceptance of beliefs to produce the largest *cut* of $content(\mathbf{B})$ so that it does not entail $\bot$. In other words, only the set of consistent beliefs $\Delta = \{\beta \in K : \bot < \beta\}$, where $<$ is the strict part of epistemic entrenchment, will be used by the filtering agent to infer the relevance of Web documents. Similar idea has been explored in developing the expectation inference relation (Gärdenfors & Makinson 1994). For instance, $\alpha \mid\sim_{\leqslant} \gamma$ iff $\gamma \in Cn(\{\alpha\} \cup \{\beta \in \Delta : \neg\alpha < \beta\})$ is developed based on $\Delta = \{\beta : \bot < \beta\}$. So, with reference to the above example, after applying $\mathbf{B}^{\star}(penguin(tweety), 0.8)$, the agent should only make use of the following set of beliefs for reasoning:

$$\Delta = \{\forall x[penguin(x) \rightarrow bird(x)],$$
$$penguin(tweety),$$
$$bird(tweety),$$
$$\forall x[penguin(x) \rightarrow \neg fly(x)],$$
$$\neg fly(tweety)\}$$

Therefore, the agent can conclude that the user is interested in information objects about *non-flying tweety*. The above reasoning process can also be explained based on *nontrivial possibilistic deduction* ($|\sim_P$) (Dubois, Lang, & Parade 1993; Dubois, Lang, & Prade 1994). In possibilistic logic, the inconsistency degree $Incons(F)$ of a possibilistic knowledge base F is defined as the least certain formula involved in the strongest contradiction of F. Moreover, nontrivial possibilistic deduction is defined as: $F \mid\sim_P (\alpha, m)$ iff $F \models (\alpha, m)$ and $m > Incons(F)$. If the entrenchment rank of a formula α is taken as the certainty m of a possibilistic formula, $Incons(K) = 0.4$ with reference to the above example. By employing possibilistic resolution, $K \models (fly(tweety), 0.4)$ and $K \models (\neg fly(tweety), 0.7)$ can be obtained, where $\models$ is possibilistic entailment. Since the certainty degree m of $\neg fly(tweety)$ equals 0.7 and is greater than $Incons(K)$, $K \mid\sim_P \neg fly(tweety)$. Nevertheless, $fly(tweety)$ can not be deduced from K based on ($|\sim_P$) because $m = 0.4 \not> Incons(K)$. However, further investigation is required to apply possibilistic based inference to the matching components of adaptive information filtering agents.

Conclusions

The AGM belief revision paradigm offers a powerful and rigorous foundation to model the changes of an agent's beliefs. The maxi-adjustment strategy, which follows the AGM rationale of consistent and minimal belief changes, provides a robust and effective computational mechanism for the development of the filtering agents' learning components. As semantic relationships among information items can be reasoned about via the maxi-adjustment method, less human intervention may be required during the agents' reinforcement learning processes. This opens the door to better learning autonomy in adaptive information filtering agents. The technical feasibility of applying the maxi-adjustment method to adaptive information filtering agents has been examined. However, quantitative evaluation of the effectiveness of these agents needs to be conducted to verify the advantages of applying such a framework to construct the learning mechanisms of these agents.

Acknowledgments The work reported in this paper has been funded in part by the Cooperative Research Centres Program through the Department of the Prime Minister and Cabinet of Australia.

References

[Alchourrón, Gärdenfors, & Makinson 1985] Alchourrón, C.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: partial meet contraction and revision functions. *Journal of Symbolic Logic* 50:510–530.

[Allan 1996] Allan, J. 1996. Incremental relevance feedback for information filtering. In *Proceedings of the 19th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, Filtering, 270–278.

[Armstrong *et al.* 1995] Armstrong, R.; Freitag, D.; Joachims, T.; and Mitchell, T. 1995. Webwatcher: A

learning apprentice for the world wide web. In *AAAI Spring Symposium on Information Gathering*, 6–12.

[Balabanovic 1997] Balabanovic, M. 1997. An adaptive web page recommendation service. In Johnson, W. L., and Hayes-Roth, B., eds., *Proceedings of the First International Conference on Autonomous Agents (Agents'97)*, 378–385. New York: ACM Press.

[Billsus & Pazzani 1999] Billsus, D., and Pazzani, M. 1999. A personal news agent that talks, learns and explains. In *Proceedings of the Third International Conference on Autonomous Agents (Agents'99)*, 268–275. Seattle, WA: ACM Press.

[Bruza & Huibers 1994] Bruza, P., and Huibers, T. 1994. Investigating Aboutness Axioms Using Information Fields. In Croft, W., and Rijsbergen, C. v., eds., *Proceedings of the 17th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, 112–121. Dublin, Ireland: Springer-Verlag.

[Buckley, Salton, & Allan 1994] Buckley, C.; Salton, G.; and Allan, J. 1994. The effect of adding relevance information in a relevance feedback environment. In *Proceedings of the Seventeenth Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, Routing, 292–300.

[Chiaramella & Chevallet 1992] Chiaramella, Y., and Chevallet, J. P. 1992. About retrieval models and logic. *The Computer Journal* 35(3):233–242.

[Dubois, Lang, & Parade 1993] Dubois, D.; Lang, J.; and Parade, H. 1993. Possibilistic Logic. In Gabbay, D. M.; Hogger, C. J.; Robinson, J. A.; and Nute, D., eds., *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume 3. Oxford: Oxford University Press. 439–513.

[Dubois, Lang, & Prade 1994] Dubois, D.; Lang, J.; and Prade, H. 1994. Automated Reasoning using Possibilistic Logic: Semantics, Belief Revision, and Variable Certainty Weights. *IEEE Transactions on Knowledge and Data Engineering* 6(1):64–71.

[Gärdenfors & Makinson 1988] Gärdenfors, P., and Makinson, D. 1988. Revisions of knowledge systems using epistemic entrenchment. In Vardi, M. Y., ed., *Proceedings of the Second Conference on Theoretical Aspects of Reasoning About Knowledge*, 83–95. San Francisco, CA: Morgan Kaufmann Inc.

[Gärdenfors & Makinson 1994] Gärdenfors, P., and Makinson, D. 1994. Nonmonotonic inference based on expectations. *Artificial Intelligence* 65(2):197–245.

[Gärdenfors 1988] Gärdenfors, P. 1988. *Knowledge in flux: modeling the dynamics of epistemic states*. Cambridge, Massachusetts: The MIT Press.

[Gärdenfors 1992] Gärdenfors, P. 1992. Belief revision: An introduction. In Gärdenfors, P., ed., *Belief Revision*. Cambridge, UK: Cambridge University Press. 1–28.

[Hunter 1995] Hunter, A. 1995. Using default logic in information retrieval. In Froidevaux, C., and Kohlas, J., eds., *Symbolic and Quantitative Approaches to Uncertainty*, volume 946 of *Lecture Notes in Computer Science*, 235–242.

[Kindo *et al.* 1997] Kindo, T.; Yoshida, H.; Morimoto, T.; and Watanabe, T. 1997. Adaptive personal information filtering system that organizes personal profiles automatically. In Pollack, M. E., ed., *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence*, 716–721. San Francisco, CA: Morgan Kaufmann publishers Inc.

[Lalmas 1998] Lalmas, M. 1998. Logical models in information retrieval: Introduction and overview. *Information Processing & Management* 34(1):19–33.

[Langley 1996] Langley, P. 1996. *Elements of Machine Learning*. San Francisco, CA: Morgan Kaufmann Publishers.

[Lau, Hofstede, & Bruza 1999] Lau, R.; Hofstede, A. H. M.; and Bruza, P. D. 1999. A Study of Belief Revision in the Context of Adaptive Information Filtering. In *Proceedings of the Fifth International Computer Science Conference (ICSC'99)*, volume 1749 of *Lecture Notes in Computer Science*, 1–10. Berlin: Springer.

[Moukas & Maes 1998] Moukas, A., and Maes, P. 1998. Amalthaea: An evolving information filtering and discovery system for the WWW. *Journal of Autonomous Agents and Multi-Agent Systems* 1(1):59–88.

[Nie, Brisebois, & Lepage 1995] Nie, J.; Brisebois, M.; and Lepage, F. 1995. Information retrieval as counterfactual. *The Computer Journal* 38(8):643–657.

[Pazzani, Muramatsu, & Billsus 1996] Pazzani, M.; Muramatsu, J.; and Billsus, D. 1996. Syskill and Webert: Identifying interesting web sites. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence and the Eighth Innovative Applications of Artificial Intelligence Conference*, 54–61. Menlo Park: AAAI Press / MIT Press.

[Salton & Buckley 1990] Salton, G., and Buckley, C. 1990. Improving retrieval performance by relevance feedback. *Journal of American Society for Information Science* 41(4):288–297.

[Salton & McGill 1983] Salton, G., and McGill, M. 1983. *Introduction to Modern Information Retrieval*. New York: McGraw-Hill.

[Salton 1989] Salton, G. 1989. *Automatic Text Processing–The Transformation, Analysis, and Retrieval of Information by Computer*. Reading, Massachusetts: Addison-Wesley.

[Williams 1995] Williams, M.-A. 1995. Iterated theory base change: A computational model. In Mellish, C. S., ed., *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, 1541–1547. San Francisco, CA: Morgan Kaufmann Publishers Inc.

[Williams 1996] Williams, M.-A. 1996. Towards a practical approach to belief revision: Reason-based change.

In Aiello, L. C.; Doyle, J.; and Shapiro, S., eds., *KR'96: Principles of Knowledge Representation and Reasoning*, 412–420. San Francisco, CA: Morgan Kaufmann Publishers Inc.

[Williams 1997a] Williams, M.-A. 1997a. Anytime belief revision. In Pollack, M. E., ed., *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence*, 74–79. San Francisco, CA: Morgan Kaufmann Publishers Inc.

[Williams 1997b] Williams, M.-A. 1997b. Implementing Belief Revision. In Antoniou, G., ed., *Nonmonotonic Reasoning*. Cambridge, Massachusetts: The MIT Press. 197–211.

[Williams 1998] Williams, M.-A. 1998. Applications of belief revision. *Lecture Notes in Computer Science* 1472:287–316.